AMENDMENTS TO THE CLAIMS

1. (Currently amended) A method for managing a data storage system that includes primary and secondary storage subsystems, including respective first and second non-volatile storage media, the method comprising:

maintaining a bitmap record on the secondary storage subsystem, which is predictive of locations to which data are to be written on the primary storage subsystem by a host processor, the record including a designation of locations to which the host is expected to write in the near future;

receiving at the primary storage subsystem, from the host processor, the data to be written to a specified location on the first non-volatile storage media;

~~if~~ making a determination that the specified location is not included in the record, and responsively to the determination sending a message from the primary storage subsystem to the secondary storage subsystem so as to cause the secondary storage subsystem to update the record,

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written the data and is expected to write the data in a subsequent write operation, and to set a number of predicted bits in the record corresponding to the one or more further locations,

wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery;

signaling the host processor that the data have been stored in the data storage system responsively to receiving the data and, ~~if~~ upon having made the determination that the specified location was not included in the record, responsively to receiving an acknowledgment at the primary storage subsystem from the

secondary storage subsystem indicating that the record has been updated; and

storing the data in the specified location on both the first and second non-volatile storage media.

2. (Original) The method according to claim 1, wherein sending the message comprises copying the data synchronously from the primary storage subsystem to the secondary storage subsystem.

3. (Currently amended) The method according to claim 2, wherein storing the data comprises, ~~if~~ upon making the determination that the specified location is included in the record, copying the data from the primary storage subsystem to the secondary storage subsystem asynchronously, without updating the record with respect to the specified location.

4. (Original) The method according to claim 3, wherein copying the data comprises transmitting the data between mutually-remote sites over a communication link between the sites.

5. (Currently amended) The method according to claim 3, wherein maintaining the record comprises maintaining a copy of the record on the primary storage subsystem, and wherein signaling the host processor comprises, ~~if~~ upon making the determination that the specified location is included in the record, indicating to the host processor that the data have been stored without waiting to receive the acknowledgment from the secondary storage subsystem.

6. (Original) The method according to claim 1, wherein copying the data comprises creating a mirror on the secondary storage subsystem of the data received by the primary storage subsystem.

7. (Original) The method according to claim 6, and comprising, upon occurrence of a failure in the primary storage subsystem, configuring the secondary storage

subsystem to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

8. (Original) The method according to claim 6, and comprising, upon recovery of the system from a failure of the primary storage subsystem, conveying, responsively to the record, a portion of the data from the secondary storage subsystem to the primary storage subsystem for storage on the primary storage subsystem.

9. (Original) The method according to claim 1, wherein maintaining and updating the record comprise marking respective bits in a bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

10. (Currently amended) A method for managing a data storage system that includes primary and secondary storage subsystems, including respective first and second non-volatile storage media, the method comprising:

maintaining a record on the secondary storage subsystem, which is predictive of locations to which data are to be written on the primary storage subsystem by a host processor, wherein maintaining the record comprises maintaining a copy of the record on the primary storage subsystem;

receiving at the primary storage subsystem, from the host processor, the data to be written to a specified location on the first non-volatile storage media;

~~if~~ making a determination that the specified location is not included in the record, and responsively to the determination sending a message from the primary storage subsystem to the secondary storage subsystem so as to cause the secondary storage subsystem to update the record, wherein sending the message comprises deciding at the primary storage subsystem to send the message responsively to the copy of the record, and

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written the data and is expected to write the data in a subsequent write operation, and to set a number of predicted bits in the record corresponding to the one or more further locations,

wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery;

signaling the host processor that the data have been stored in the data storage system responsively to receiving the data and, ~~if~~ upon having made the determination that the specified location was not included in the record, responsively to receiving an acknowledgment at the primary storage subsystem from the secondary storage subsystem indicating that the record has been updated; and

storing the data in the specified location on both the first and second non-volatile storage media.

11. (Original) The method according to claim 10, wherein sending the message comprises modifying both the record and the copy of the record responsively to the specified location.

12. (Original) The method according to claim 11, wherein modifying both the record and the copy of the record comprises adding a plurality of locations, including the specified location, to both the record and the copy of the record.

13. (Original) The method according to claim 10, wherein maintaining the copy of the record comprises selecting one or more locations, other than the specified location, to be removed from the record, and instructing the secondary storage subsystem to remove the one or more

locations from the record, so as to limit a size of the record.

14. (Original) The method according to claim 13, wherein storing the data comprises copying the data to be stored in the one or more locations from the primary storage subsystem to the secondary storage subsystem, and wherein selecting the one or more locations comprises receiving a return message from the secondary storage subsystem indicating that the secondary storage subsystem has received the copied data, and selecting the one or more locations to be removed from the record responsively to the return message.

15. (Original) The method according to claim 13, wherein selecting the one or more locations comprises identifying the locations at which the first and second non-volatile storage media contain substantially identical data, and selecting for removal one of the identified locations that was least-recently added to the record.

16. (Original) The method according to claim 13, wherein sending the message comprises adding one or more entries to both the record and the copy of the record responsively to the specified location, and grouping the entries added to the copy of the record and the record in generations according to an order of adding the entries to the records, and wherein selecting the one or more locations comprises determining at the primary subsystem that all the entries in one of the generations may be removed from the record.

17. (Original) The method according to claim 13, wherein instructing the secondary storage subsystem to remove the one or more locations comprises appending an instruction to the message sent from the primary storage subsystem to the secondary storage subsystem.

18. (Canceled)

19. (Currently amended) The method according to ~~claim 18~~ claim 1, wherein the one or more further locations comprise a predetermined number of consecutive locations in proximity to the specified location.

20. (Currently amended) The method according to ~~claim 18~~ claim 1, wherein maintaining the record comprises recording the locations to which the data are written using an object-based storage technique, and wherein the one or more further locations are chosen based on a logical connection between storage objects.

21. (Currently amended) A data storage system, comprising:

a primary storage subsystem, which comprises first non-volatile storage media; and

a secondary storage subsystem, which comprises second non-volatile storage media, and which is arranged to maintain a bitmap record that is predictive of locations to which data are to be written on the primary storage subsystem by a host processor, the record including a designation of locations to which the host is expected to write in the near future,

wherein the primary storage subsystem is arranged to receive the data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, which is arranged to store the data in the specified location on the second non-volatile storage media, and

wherein the primary storage subsystem is further arranged, upon receiving from the host processor the data to be written to a specified location on the first non-volatile storage media, ~~if~~ to make a determination that the specified location is not included in the record, and responsively to the determination to send a message to the secondary storage subsystem so as to cause the

secondary storage subsystem to update the record and to return an acknowledgment to the primary storage subsystem indicating that the record has been updated,

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written the data and is expected to write the data in a subsequent write operation, and to set a number of predicted bits in the record corresponding to the one or more further locations,

wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery, and

wherein the primary storage subsystem is further arranged to signal the host processor that the data have been stored in the data storage system responsively to receiving the data and, ~~if~~ upon having made the determination that the specified location was not included in the record, responsively to receiving the acknowledgment from the secondary storage subsystem.

22. (Original) The system according to claim 21, wherein the message sent to the secondary storage subsystem comprises the data, which are copied synchronously from the primary storage subsystem to the secondary storage subsystem.

23. (Currently amended) The system according to claim 22, wherein the primary storage subsystem is arranged, ~~if~~ upon making the determination that the specified location is included in the record, to copy the data from the primary storage subsystem to the secondary storage subsystem asynchronously, without causing the secondary storage subsystem to update the record with respect to the specified location.

24. (Original) The system according to claim 23, wherein the first and second non-volatile storage media are

located at mutually-remote sites, and wherein at least one of the primary and secondary storage subsystems is arranged to transmit the data over a communication link between the sites.

25. (Currently amended) The system according to claim 23, wherein the primary storage subsystem is arranged to maintain a copy of the record, and, ~~if~~ upon making the determination that the specified location is included in the record, to signal to the host processor that the data have been stored without waiting to receive the acknowledgment from the secondary storage subsystem.

26. (Original) The system according to claim 21, wherein the secondary storage subsystem is arranged to mirror the data held by the primary storage subsystem.

27. (Original) The system according to claim 26, wherein upon occurrence of a failure in the primary storage subsystem, the secondary storage subsystem is configurable to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

28. (Original) The system according to claim 26, wherein upon recovery of the system from a failure of the primary storage subsystem, the secondary storage subsystem is arranged to convey, responsively to the record, a portion of the data from the second non-volatile storage media to the primary storage subsystem for storage on the first non-volatile storage media.

29. (Original) The system according to claim 21, wherein the record comprises a bitmap, and wherein the secondary storage subsystem is arranged to mark respective bits in the bitmap corresponding to the locations to which the data are to be written by the host processor.

30. (Currently amended) A data storage system, comprising:

a primary storage subsystem, which comprises first non-volatile storage media; and

a secondary storage subsystem, which comprises second non-volatile storage media, and which is arranged to maintain a record that is predictive of locations to which data are to be written on the primary storage subsystem by a host processor,

wherein the primary storage subsystem is arranged to receive the data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, which is arranged to store the data in the specified location on the second non-volatile storage media, and

wherein the primary storage subsystem is further arranged, upon receiving from the host processor the data to be written to a specified location on the first non-volatile storage media, ~~if~~ to make a determination that the specified location is not included in the record, and responsively to the determination to send a message to the secondary storage subsystem so as to cause the secondary storage subsystem to update the record and to return an acknowledgment to the primary storage subsystem indicating that the record has been updated, wherein the primary storage subsystem is arranged to maintain a copy of the record, and to determine whether to send the message responsively to the copy of the record,

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written the data and is expected to write the data in a subsequent write operation, and to set a number of predicted bits in the record corresponding to the one or more further locations,

wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery, and

wherein the primary storage subsystem is further arranged to signal the host processor that the data have been stored in the data storage system responsively to receiving the data and, if the specified location was not included in the record, responsively to receiving the acknowledgment from the secondary storage subsystem.

31. (Original) The system according to claim 30, wherein the primary and secondary storage subsystems are arranged to update the copy of the record and the record, respectively, responsively to the specified location.

32. (Original) The system according to claim 31, wherein the primary and secondary storage subsystems are arranged to update the copy of the record and the record by adding a plurality of locations, including the specified location, to both the first and second records.

33. (Original) The system according to claim 30, wherein the primary storage subsystem is arranged to select one or more locations, other than the specified location, to be removed from the copy of the record, and to instruct the secondary storage subsystem to remove the one or more locations from the record, so as to limit a size of the record.

34. (Original) The system according to claim 33, wherein the secondary storage subsystem is arranged to send a return message to the primary storage subsystem, indicating that the secondary storage subsystem has received the copied data, and wherein the primary storage subsystem is arranged to select the one or more locations to be removed from the record responsively to receiving the return message.

35. (Original) The system according to claim 33, wherein the primary storage subsystem is arranged to identify the

locations at which the first and second non-volatile storage media contain substantially identical data, and to select for removal one of the identified locations that was least-recently added to the record.

36. (Original) The system according to claim 33, wherein the primary and secondary storage subsystems are arranged to respectively add one or more entries to both the copy of the record and the record responsively to the specified location, and to group the entries added to the first and second records in generations according to an order of adding the entries to the records, and wherein the primary storage subsystem is arranged to determine that all the entries in one of the generations may be removed from the record, and to instruct the secondary storage subsystem to remove all the entries in the one of the generations from the record.

37. (Original) The system according to claim 33, wherein the primary storage subsystem is arranged to append an instruction to the message sent to the secondary storage subsystem, so as to instruct the secondary storage subsystem to remove the one or more locations from the record.

38. (Canceled)

39. (Currently amended) The system according to ~~claim 38~~ claim 21, wherein the one or more further locations comprise a predetermined number of consecutive locations in proximity to the specified location.

40. (Currently amended) The system according to ~~claim 38~~ claim 21, wherein the secondary storage subsystem is arranged to maintain the record using an object-based storage technique, and to predict the one or more further locations based on a logical connection between storage objects.

41. (Currently amended) A computer software product for use in a data storage system including primary and secondary storage subsystems, which include respective first and second control units and respective first and second non-volatile storage media, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the first and second control units, cause the first control unit to receive data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, and cause the second control unit to maintain a bitmap record that is predictive of locations to which the data are to be written on the primary storage subsystem by the host processor, the record including a designation of locations to which the host is expected to write in the near future, and to store the data copied to the second storage subsystem in the specified location on the second non-volatile storage media,

wherein the instructions further cause the first control unit, ~~if~~ to make a determination that the specified location is not included in the record, and responsively to the determination to send a message to the secondary storage subsystem so as to cause the second control unit to update the record and to return an acknowledgment to the primary storage subsystem, and cause the first control unit to signal the host processor that the data have been stored in the data storage product responsively to receiving the data and, ~~if~~ upon having made the determination that the specified location was not included in the record, responsively to receiving the acknowledgment from the second control unit,

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written

the data and is expected to write the data in a subsequent write operation, and to set a number of predicted bits in the record corresponding to the one or more further locations,

wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery.

42. (Original) The product according to claim 41, wherein the message sent to the secondary storage subsystem comprises the data, which are copied synchronously from the primary storage subsystem to the secondary storage subsystem.

43. (Currently amended) The product according to claim 42, wherein the instructions cause the first control unit, ~~if~~ upon making the determination that the specified location is included in the record, to copy the data from the primary storage subsystem to the secondary storage subsystem asynchronously, without causing the second control unit to update the record with respect to the specified location.

44. (Original) The product according to claim 43, wherein the first and second non-volatile storage media are located at mutually-remote sites, and wherein the instructions cause at least one of the first and second control units to transmit the data over a communication link between the sites.

45. (Currently amended) The product according to claim 43, wherein the instructions cause the first control unit to maintain a copy of the record, and, ~~if~~ upon making the determination that the specified location is included in the record, to signal to the host processor that the data have been stored without waiting to receive the acknowledgment from the second control unit.

46. (Original) The product according to claim 41, wherein the instructions cause the first and second

control units to mirror the data held by the primary storage subsystem on the secondary storage subsystem.

47. (Original) The product according to claim 46, wherein the instructions cause the secondary storage subsystem, upon occurrence of a failure in the primary storage subsystem, to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

48. (Original) The product according to claim 46, wherein upon recovery of the system from a failure of the primary storage subsystem, the instructions cause the second control unit to convey, responsively to the record, a portion of the data from the second non-volatile storage media to the primary storage subsystem for storage on the first non-volatile storage media.

49. (Original) The product according to claim 41, wherein the record comprises a bitmap, and wherein the instructions cause the second control unit to mark respective bits in the bitmap corresponding to the locations to which the data are to be written by the host processor.

50. (Currently amended) A computer software product for use in a data storage system including primary and secondary storage subsystems, which include respective first and second control units and respective first and second non-volatile storage media, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the first and second control units, cause the first control unit to receive data from a host processor for writing to a specified location, and to store the data in the specified location on the first non-volatile storage media while copying the data to the second storage subsystem, and cause the second control unit to maintain a record that is predictive of locations to which the

data are to be written on the primary storage subsystem by the host processor, and to store the data copied to the second storage subsystem in the specified location on the second non-volatile storage media, wherein the instructions cause the first control unit to maintain a copy of the record, and to determine whether to send the message responsively to the copy of the record, and

wherein the instructions further cause the first control unit, ~~if~~ to make a determination that the specified location is not included in the record, and responsively to the determination to send a message to the secondary storage subsystem so as to cause the second control unit to update the record and to return an acknowledgment to the primary storage subsystem, and cause the first control unit to signal the host processor that the data have been stored in the data storage product responsively to receiving the data and, ~~if~~ upon having made the determination that the specified location was not included in the record, responsively to receiving the acknowledgment from the second control unit,

wherein sending the message causes the secondary storage subsystem to predict one or more further locations to which the host processor has not yet written the data and is expected to write the data in a subsequent write operation, and to set a number of predicted bits in the record corresponding to the one or more further locations,

wherein the number of the predicted bits is chosen so as to achieve a desired balance between low average latency and rapid failure recovery.

51. (Original) The product according to claim 50, wherein the instructions cause the first and second control units to update the copy of the record and the record, respectively, responsively to the specified location.

52. (Original) The product according to claim 51, wherein the instructions cause the first and second control units to update the copy of the record and the record by adding a plurality of locations, including the specified location, to both the first and second records.

53. (Original) The product according to claim 50, wherein the instructions cause the first control unit to select one or more locations, other than the specified location, to be removed from the copy of the record, and to instruct the second control unit to remove the one or more locations from the record, so as to limit a size of the record.

54. (Original) The product according to claim 53, wherein the instructions cause the second control unit to send a return message to the primary storage subsystem, indicating that the secondary storage subsystem has received the copied data, and wherein the instructions cause the first control unit to select the one or more locations to be removed from the record responsively to receiving the return message.

55. (Original) The product according to claim 53, wherein the instructions cause the first control unit to identify the locations at which the first and second non-volatile storage media contain substantially identical data, and to select for removal one of the identified locations that was least-recently added to the record.

56. (Original) The product according to claim 53, wherein the instructions cause the first and second control units to respectively add one or more entries to both the copy of the record and the record responsively to the specified location, and to group the entries added to the first and second records in generations according to an order of adding the entries to the records, and wherein the instructions cause the first control unit to determine that all the entries in one of the generations

may be removed from the records, and to instruct the second control unit to remove all the entries in the one of the generations from the record.

57. (Original) The product according to claim 53, wherein the instructions cause the first control unit to append an instruction to the message sent to the secondary storage subsystem, so as to instruct the second control unit to remove the one or more locations from the record.

58. (Canceled)

59. (Currently amended) The product according to ~~claim 58~~ claim 41, wherein the one or more further locations comprise a predetermined number of consecutive locations in proximity to the specified location.

60. (Currently amended) The product according to ~~claim 58~~ claim 41, wherein the instructions cause the second control unit to maintain the record using an object-based storage technique, and to predict the one or more further locations based on a logical connection between storage objects.